Deficit Round-Robin-Based ISP Traffic Control Scheme Enabling Excess Bandwidth Allocation in Shared Access Networks

Kyeong Soo Kim, *Member, IEEE*

arXiv:1403.5712v2 [cs.NI] 25 Mar 2014

***Abstract*—In shared access, shaping subscriber traffic based on token bucket by ISPs results in waste of network resources when there are few active subscribers, because it cannot allocate excess bandwidth in the long term. To address it, a traffic control scheme based on core-stateless fair queueing (CSFQ) and token bucket meters (TBMs) has been recently proposed, which can allocate excess bandwidth among active subscribers proportional to their token generation rates. The use of FIFO queue for all packets, however, degrades the short-term performance of conformant traffic due to the presence of non-conformant packets already in the queue. Also, the rate estimation based on exponential averaging makes it difficult to react to rapid changes in traffic conditions. In this letter we propose a new traffic control scheme based on deficit round-robin (DRR) and TBMs to guarantee the quality of service of conformant packets in all time scales while allocating excess bandwidth among active subscribers proportional to their token generation rates, whose advantages over the CSFQ-based scheme are demonstrated through simulation results.**

***Index Terms*—Access, Internet service provider (ISP), traffic shaping, fair queueing, deficit round-robin (DRR), quality of service (QoS).**

I. Introduction

RECENTLY, an Internet service provider (ISP) traffic control scheme based on core-stateless fair queueing (CSFQ) [1] and token bucket meters (TBMs) has been proposed to allocate excess bandwidth among active subscribers in a fair and efficient way, while not compromising the service contracts specified by the original token bucket filter (TBF) for conformant subscribers [2]. Through the use of a common first in, first out (FIFO) queue for both conformant and non-conformant packets, it can preserve packet sequence; handling conformant and non-conformant packets differently at their arrivals, it gives priority to the former, while allocating excess bandwidth to the latter proportional to their token generation rates. In this way, it addresses the critical issue in traffic shaping based on the original TBF that the excess bandwidth, resulting from the inactivity of some subscribers, cannot be allocated to other active subscribers in the long term.

The use of a common FIFO queue in the CSFQ-based scheme, however, degrades the short-term performance of conformant traffic due to the presence of non-conformant packets already in the queue. Also, the rate estimation based on exponential averaging makes it difficult to quickly react to rapid changes in traffic conditions. These may compromise the quality of service (QoS) for conformant traffic compared to that with the original TBF in the short term, especially for bursty, elastic traffic like that of transmission control protocol (TCP).

K. S. Kim is with the College of Engineering, Swansea University, Swansea, SA2 8PP, Wales United Kingdom (e-mail: k.s.kim@swansea.ac.uk).

In order to address the issues resulting from the use of common FIFO queue and rate estimation based on exponential averaging in the CSFQ-based scheme, in this letter we propose a new ISP traffic control scheme based on deficit round-robin (DRR) [3], which ideally combines the advantages of both the original TBF (i.e., passing short, bursty conformant traffic without shaping) and the weighted fair queueing (WFQ) (i.e., allocating excess bandwidth among active flows proportional to their weights).

II. Excess Bandwidth Allocation based on DRR

The proposed ISP traffic control scheme is to meet the following requirements in allocating excess bandwidth [2]: First, the allocation of excess bandwidth should not compromise the QoS of traffic conformant to service contracts based on the *original token bucket* algorithm; secondly, excess bandwidth should be allocated among active subscribers proportional to their negotiated long-term average rates, i.e., token generation rates.

The first requirement implies that conformant packets should have priority over non-conformant ones in queueing and scheduling. The second requirement can be stated formally as follows [2]: Let $C_{ex}(t)$ and $A(t)$ be *excess bandwidth* and *total arrival rate of non-conformant packets* at time t for a shared access network with N subscribers, i.e.,

$$C_{ex}(t) \triangleq C - r_c(t), \tag{1}$$

and

$$A(t) \triangleq \sum_{i=1}^{N} r_{nc,i}(t), \tag{2}$$

where C is the capacity of the access link, $r_c(t)$ the arrival rate of conformant packets for all subscribers, and $r_{nc,i}(t)$ the arrival rate of non-conformant packets for the ith subscriber, respectively. If $A(t) > C_{ex}(t)$, the *normalized* fair rate $\alpha(t)$ is a unique solution to

$$C_{ex}(t) = \sum_{i=1}^{N} w_i \min(\alpha(t),\ r_{nc,i}(t)/w_i), \tag{3}$$

where w_i is the weight for the ith subscriber, which is proportional to the token generation rate; otherwise, $\alpha(t)$ is set to $max_i\,(r_{nc,i}(t)/w_i)$ [1].

Fig. 1 shows two ISP traffic control schemes enabling proportional allocation of excess bandwidth, i.e., the CSFQ-based scheme [2] and the proposed DRR-based one. Unlike the CSFQ-based scheme, the proposed DRR-based scheme uses per-subscriber queues, which separate traffic from different subscribers, but can preserve packet sequence using a common per-subscriber queue for both conformant and

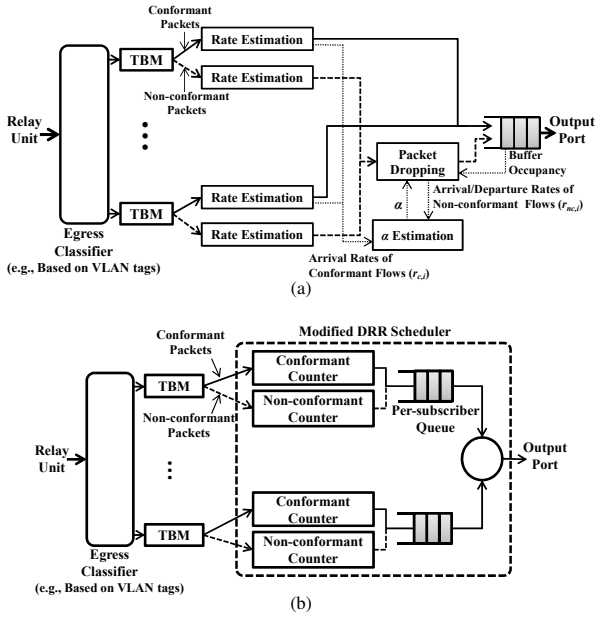

Fig. 1. ISP traffic control schemes enabling proportional allocation of excess bandwidth based on (a) CSFQ [2] and (b) DRR (deficit counters are not shown).

non-conformant packets belonging to the same subscriber. Algorithms 1 and 2 show pseudocode of enqueueing and dequeueing procedures of the proposed scheme, where DC_i, CC_i, NC_i and Q_i are deficit, conformant byte, non-conformant byte counters and quantum for the ith subscriber, respectively. $ConformantList$ and $NonconformantList$ are lists of active queues having conformant and non-conformant packets. The two additional counters (i.e., CC_i and NC_i) are used to keep track of the number of bytes for conformant and non-conformant packets in a queue, which function as *logically separate* queues within a common per-subscriber queue for sequence preserving; to give priority to conformant packets in queueing, when a newly arrived conformant packet is to be discarded due to buffer overflow, NC_i is decreased instead, while CC_i is increased, which emulates preemptive queueing. Likewise, the two lists of active queues are used to give conformant packets priority in scheduling by checking $ConformantList$ first during the dequeueing procedure; as described in Algorithm 2, conformant packets are first scheduled in a round-robin manner (i.e., without taking into account deficit counters), while non-conformant packets, after serving all conformant packets in the queues, are scheduled based on DRR for proportional allocation of excess bandwidth.

Note that the proposed scheme does not use the rate estimation based on exponential averaging that makes it difficult for the CSFQ-based scheme to react promptly to rapid changes in traffic conditions and interact with TCP flows. Also note that the proposed DRR-based scheme, whose complexity is $\mathcal{O}(1)$, has an advantage in complexity over the CSFQ-based scheme which corresponds to the extreme case of CSFQ islands, i.e., the node itself is an island, where both the functionalities of edge and core routers of CSFQ reside in the same node.

On receiving a packet p for the ith subscriber:
$PacketSize \leftarrow Size(p)$ */* in byte */*
$Conformed \leftarrow Meter(p)$
$Dropped \leftarrow Enqueue(i,\ p)$
if $Dropped == TRUE$
  if $Conformed == TRUE$ */* packet swapping */*
    $CC_i \leftarrow CC_i + PacketSize$
    $NC_i \leftarrow NC_i - PacketSize$
else
  if $Conformed == TRUE$
    $CC_i \leftarrow CC_i + PacketSize$
    if $i \notin ConformantList$
      Add i to the end of $ConformantList$
  else
    $NC_i \leftarrow NC_i + PacketSize$
    if $i \notin NonconformantList$
      Add i to the end of $NonconformantList$
      $DC_i \leftarrow 0$ */* reset counter */*

Algorithm 1: Enqueueing procedure.

On receiving a packet when all queues are empty or at the end of packet transmission:
while $ConformantList$ *is not empty*
  Remove head of $ConformantList$, say the ith subscriber
  $PacketSize \leftarrow Size(Head(Queue_i))$
  if $CC_i \geq PacketSize$
    $CC_i \leftarrow CC_i - PacketSize$
    $Send(Dequeue(Queue_i))$
    if $Queue_i$ *is not empty AND*
    $CC_i \geq Size(Head(Queue_i))$
      Add i to the end of $ConformantList$
    Exit */* exit here */*
while *NonconformantList is not empty*
  Remove head of $NonconformantList$, say the ith subscriber
  $PacketSize \leftarrow Size(Head(Queue_i))$
  if $Continued == TRUE$ */* from previous TX */*
    $DC_i \leftarrow 0$
  else
    $DC_i \leftarrow Q_i$;
  if $DC_i \geq PacketSize$ *AND* $NC_i \geq PacketSize$
    $DC_i \leftarrow DC_i - PacketSize$
    $NC_i \leftarrow NC_i - PacketSize$
    $Send(Dequeue(Queue_i))$
    if $Queue_i$ *is not empty*
      $PacketSize \leftarrow Size(Head(Queue_i))$
      if $NC_i \geq PacketSize$
        if $DC_i \geq PacketSize$
          $Continued \leftarrow TRUE$
          Add i to the front of $NonconformantList$
        else
          $Continued \leftarrow FALSE$
          Add i to the end of $NonconformantList$
        Exit */* exit here */*
    $Continued \leftarrow FALSE$
    $DC_i \leftarrow 0$
    Exit */* exit here */*

Algorithm 2: Dequeueing procedure.

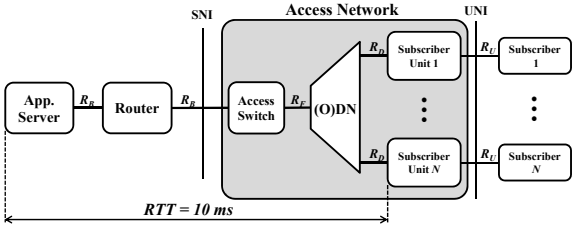

Fig. 2. A simulation model for a shared access network.

III. Simulation Results

We compared the proposed scheme with the original TBF[1] and the CSFQ-based one with buffer-based amendment using the simulation model described in [2], which is shown in Fig. 2. Subscribers are connected through 100-Mb/s user-network interfaces (UNIs) to shared access with the same feeder and distribution rates of 100-Mb/s, each of which receives packet streams from user datagram protocol (UDP) or TCP sources in the application server. The backbone rate (i.e., R_B) and the end-to-end round-trip time are set to 10 Gb/s and 10 ms.

In the first experiment, 16 subscribers are divided into 4 groups, 4 subscribers per each as in [2]: For Groups 1-3, each subscriber receives a 1000-byte packet from a UDP source at every 0.5 ms, resulting in the rate of 16 Mb/s. Token generation rates are set to 2.5 Mb/s, 5 Mb/s, and 7.5 Mb/s for Groups 1, 2, and 3, respectively. Their starting times are set to 0 s, 60 s, 120 s. For Group 4, each subscriber receives packets from a greedy TCP source with token generation rate of 10 Mb/s and starting time of 180 s. Token bucket size is set to 1 MB for all subscribers, and peak rate control is not used. The size of per-subscriber queues for both the original TBF (denoted as "RR+TBF") and the proposed scheme ("DRR+TBM") is set to 1 MB (16 MB in total). For CSFQ-based scheme ("CSFQ+TBF"), the size of common FIFO queue is set to 16 MB to cope with worst-case bursts resulting from 16 token buckets. The averaging constants for the estimation of flow rates (i.e, K) and the normalized fair rate (i.e., K_α) are set to 100 ms and 200 ms, respectively; as for the buffer-based amendment, we set a threshold to 64 kB.

Fig. 3 shows flow throughput averaged over a 1-s interval from one sample run, demonstrating dynamic performance of each scheme (i.e., *how quickly* it can respond to the changes in incoming traffic and allocate excess bandwidth accordingly). As expected, the original TBF scheme cannot allocate excess bandwidth to active subscribers except for short periods of time around 60 s, 120 s and 180 s enabled by tokens in the buckets. The CSFQ-based and the proposed schemes, on the other hand, can allocate well excess bandwidth among UDP flows until 180 s when TCP flows start; of the two, the proposed scheme provides much better performance in terms of fluctuation. Due to 1-MB token buckets, there are spikes in the throughput of newly started flows at 60 s (i.e., Group 2) and 120 s (i.e., Group 3), during which the throughput of existing flows decreases temporarily. As TCP flows start at 180 s, the difference between the CSFQ-based and the proposed schemes become even clearer: As for the CSFQ-based scheme, the buffer-based amendment reduces the transient period, but at the expense of fluctuations in steady states. With the proposed scheme, there is virtually no fluctuation in TCP flow throughputs as well, but there is small increase in TCP throughput which lasts from 180 s to 197 s due to 1-MB token buckets, which is also the case for the original TBF.

Fig. 4 shows the average throughput of flows for two 40-s periods — i.e., a subperiod (60 s) minus a transient period (20 s) — with 95 percent confidence intervals from 10 repetitions, demonstrating static performance of each scheme (i.e., *how exactly* it can allocate available bandwidth among subscribers per the requirements described in Sec. II in a steady state). As shown in Fig. 4 (a), the CSFQ-based scheme suffers from the fluctuations observed in Fig. 3, while the proposed scheme allocates excess bandwidth from Group 4, which is inactive during this period, exactly per (3). Fig. 4 (b) shows that with TCP flows, the difference between the actual throughput of a flow and its fair share — indicated by dotted lines — become larger for the CSFQ-based scheme; note that during this period, because there is no excess bandwidth available, each flow should be allocated bandwidth per its token generation rate, which is why the original TBF scheme shows as good performance as the proposed one.

To further investigate transient responses of traffic control schemes in shorter time scale, we also carried out another experiment where we consider 4 subscribers with token generation rate of 10 Mb/s and token bucket size of 10 MB; subscriber 1 receives a 10-MB conformant burst from the application server, while subscribers 2-4 receive non-conformant UDP traffic with source rate of 50 Mb/s. The flow throughput is averaged over a 10ms-interval to better show the details. Fig. 5 illustrates the flow throughput before, during, and after the conformant burst for all three traffic control schemes, where we can clearly see that the proposed scheme provides the advantages of both the original TBF (i.e., passing the conformant burst without shaping and thereby any additional delay) and the WFQ (i.e., proportional allocation of excess bandwidth among active subscribers). In case of the CSFQ-based scheme, however, the beginning of the burst is delayed by 1.11 s due to the presence of non-conformant packets already in the FIFO queue. During the burst, the allocation of bandwidth is quite distorted (i.e., Subscriber 1 takes all the bandwidth) because the CSFQ-based scheme cannot respond quickly enough for traffic changes in such a short period of time. There is also delay after the burst in recovering the fair share of each subscriber.

IV. Conclusions

In this letter we have proposed a new DRR-based ISP traffic control scheme providing the advantages of both TBF (i.e., passing short, bursty conformant traffic without shaping) and WFQ (i.e., allocating excess bandwidth among active subscribers proportional to their token generation rates). Simulation results have demonstrated that the proposed scheme based on DRR can guarantee the QoS of conformant packets

[1]We assume that a round-robin (RR) scheduler is used with TBF. Note that both the CSFQ-based and the proposed schemes are based on schedulers.

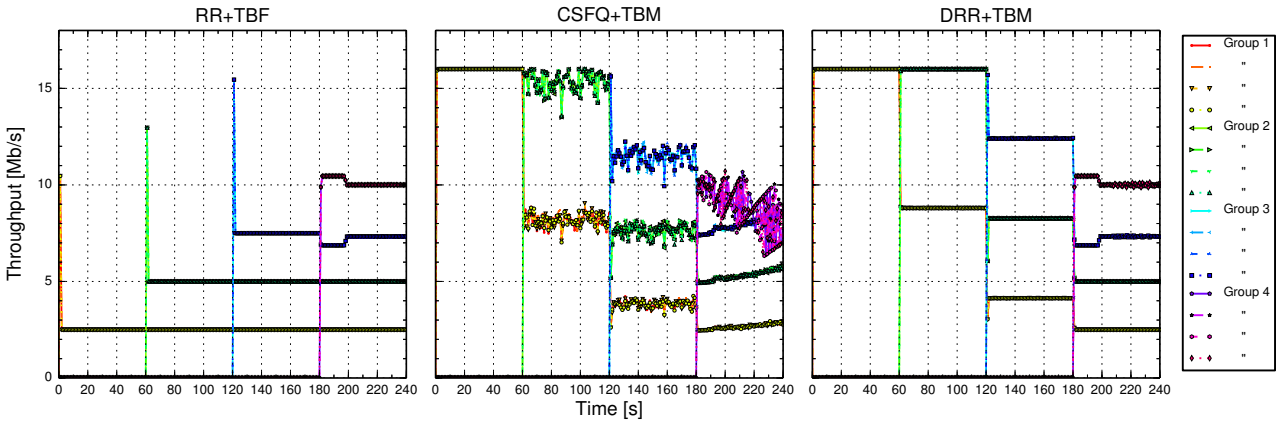

Fig. 3. Time series of throughput of flows for ISP traffic control schemes with mixed traffic.

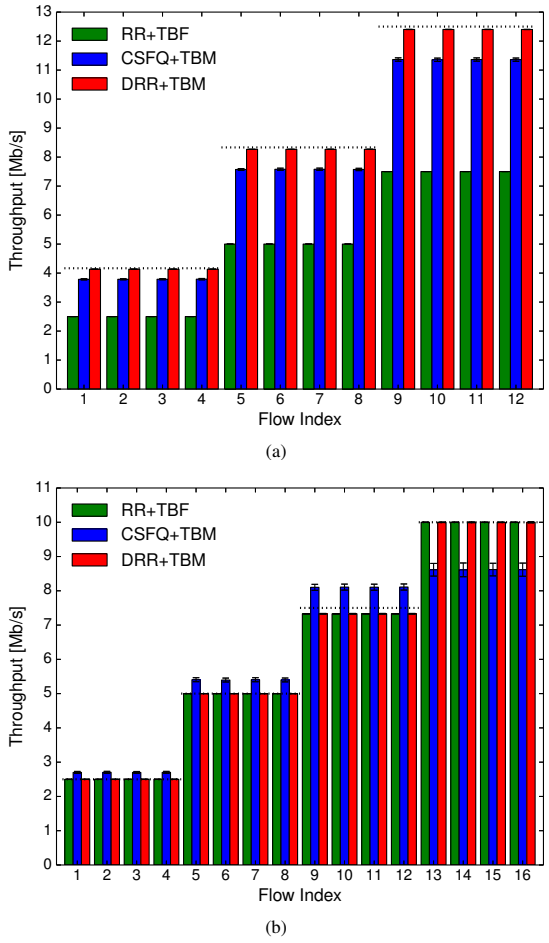

Fig. 4. Average throughput of flows with 95 percent confidence intervals for the period of (a) $140 \leq t < 180$ and (b) $200 \leq t < 240$.

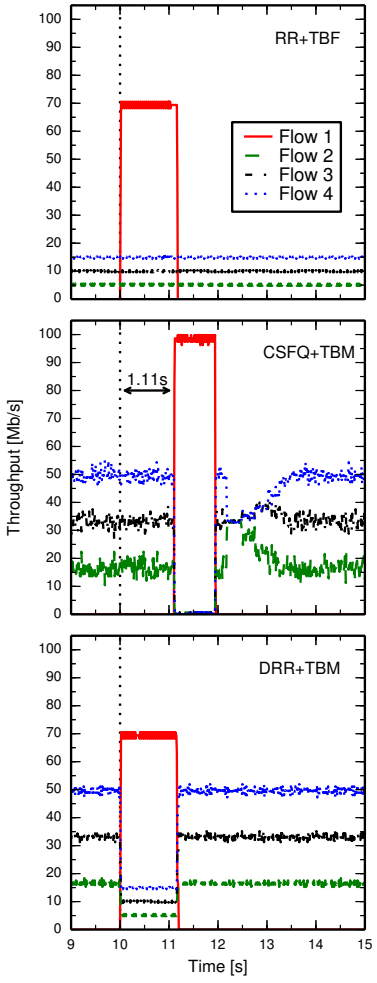

Fig. 5. Handling of a conformant burst by ISP traffic control schemes.

in all time scales while allocating excess bandwidth among active subscribers proportional to their token generation rates.

REFERENCES

[1] I. Stoica, S. Shenker, and H. Zhang, "Core-stateless fair queueing: A scalable architecture to approximate fair bandwidth allocations in high-speed networks," *IEEE/ACM Trans. Netw.*, vol. 11, no. 1, pp. 33–46, Feb. 2003.

[2] K. S. Kim, "On the excess bandwidth allocation in ISP traffic control for shared access networks," *IEEE Commun. Lett.*, no. 99, pp. 1–4, Mar. 2014. [Online]. Available: http://dx.doi.org/10.1109/LCOMM.2014.020414.132542

[3] M. Shreedhar and G. Varghese, "Efficient fair queueing using deficit round-robin," *IEEE/ACM Trans. Netw.*, vol. 4, no. 3, pp. 375–385, Jun. 1996.